arXiv:1501.01084v2 [cs.IT] 16 Feb 2015

Upper Bound on Function Computation in Directed Acyclic Networks

Cupjin Huang*, Zihan Tan* and Shenghao Yang†
*The Institute for Theoretical Computer Science (ITCS), Institute for Interdisciplinary Information Sciences,
Tsinghua University, Beijing, China
†Institute of Network Coding, The Chinese University of Hong Kong, Hong Kong, China

***Abstract*—Function computation in directed acyclic networks is considered, where a sink node wants to compute a target function with the inputs generated at multiple source nodes. The network links are error-free but capacity-limited, and the intermediate network nodes perform network coding. The target function is required to be computed with zero error. The computing rate of a network code is measured by the average number of times that the target function can be computed for one use of the network. We propose a cut-set bound on the computing rate using an equivalence relation associated with the inputs of the target function. Our bound holds for general target functions and network topologies. We also show that our bound is tight for some special cases where the computing capacity can be characterized.**

I. INTRODUCTION

We consider function computation in a directed acyclic network, where a *target function* f is intended to be calculated at a sink node, and the input symbols of the target function are generated at multiple source nodes. As a special case, network communication is just the computation of the *identity function*.[1] Network function computation naturally arises in sensor networks [1] and Internet of Things, and may find applications in big data processing.

Various models and special cases of this problem have been studied in literature (see the summarizations in [2]–[4]). We are interested in the following *network coding model* for function computation. Specifically, we assume that the network links have limited (unit) capacity and are error-free. Each source node generates multiple input symbols, and the network codes perform vector network coding by using the network multiple times.[2] An intermediate network node can transmit the output of a certain fixed function of the symbols it receives. Here all the intermediated nodes are considered with unbounded computing ability. The target function is required to be computed correctly for all possible inputs. We are interested in the *computing rate* of a network code that computes the target function, i.e., the average number of times that the target function can be computed for one use of the network. The maximum achievable computing rate is called the *computing capacity*.

When computing the identity function, the problem becomes the extensively studied network coding [5], [6], and it is known that in general linear network codes are sufficient to achieve the multicast capacity [6], [7]. For linear target functions over a finite field, a complete characterization of the computing capacity is not available for networks with one sink node. Certain necessary and sufficient conditions have been obtained such that linear network codes are sufficient to calculate a linear target function [4], [8]. But in general, linear network codes are not sufficient to achieve the computing capacity of linear target functions [9].

Networks with a single sink node are discussed in this paper, while both the target function and the network code can be non-linear. In this scenario, the computing capacity is known when the network is a multi-edge tree [2] or when the target function is the identity function. For the general case, various bounds on the computing capacity based on cut sets have been studied [2], [3]. But we find that the upper bounds claimed in [2], [3] are not valid. Specifically, the proof of [2, Theorem II.1] has an error[3]: The condition provided in the beginning of the second paragraph is not always necessary, which is illustrated by an example given in this paper. We show that the computing capacity of our example is strictly larger than the two upper bounds claimed in [2], [3].

Towards a general upper bound, we define an equivalence relation associated with the inputs of the target function (but does not depend on the network topology) and propose a cut-set bound on the computing capacity using this equivalence relation. Our bound holds for general target functions and general network topologies in the network coding model. We also show that our bound is tight when the network is a multi-edge tree or when the target function is the identity function.

In the remainder of this paper, Section II formally introduces the network computing model. The upper bound of the computing rate is given in Theorem 3, and is proved in Section IV. Section III compares with the previous results and discusses the tightness of our upper bound.

This work was supported in part by the National Basic Research Program of China Grant 2011CBA00300, 2011CBA00301, the National Natural Science Foundation of China Grant 61033001, 61361136003, 61471215. The work described in this paper was partially supported by a grant from University Grants Committee of the Hong Kong Special Administrative Region, China (Project No. AoE/E-02/08).

---

[1] A function $f : \mathcal{A} \to \mathcal{A}$ is identity if $f(x) = x$ for all $x \in \mathcal{A}$.

[2] One use of a network means the use of each link in the network at most once.

[3] No proof is provided for the upper bound (Lemma 3) in [3].

II. MAIN RESULTS

In this section, we will first introduce the network computing model. Then we will define cut sets and discuss some special cases of the function computation problem. Last we head to the main theorem about the cut-set bound for function computation.

A. Function-Computing Network Codes

Let $G = (\mathcal{V}, \mathcal{E})$ be a directed acyclic graph (DAG) with a finite vertex set $\mathcal{V}$ and an edge set $\mathcal{E}$, where multi-edges between a certain pair of nodes are allowed. A *network* over G is denoted as $\mathcal{N} = (G, S, \rho)$, where $S \subset \mathcal{V}$ is called the *source nodes* and $\rho \in \mathcal{V} \backslash S$ is called the *sink node* ρ. Let $s = |S|$, and without loss of generality (WLOG), let $S = \{1, 2, \ldots, s\}$. For an edge $e = (u, v)$, we call u the tail of e (denoted by $\text{tail}(e)$) and v the head of e (denoted by $\text{head}(e)$). Moreover, for each node $u \in \mathcal{V}$, let $\mathcal{E}_{\mathrm{i}}(u) = \{e \in \mathcal{E} : \text{head}(e) = u\}$ and $\mathcal{E}_{\mathrm{o}}(u) = \{e \in \mathcal{E} : \text{tail}(e) = u)$ be the set of incoming edges and the set of outgoing edges of u, respectively. Fix an order of the vertex set $\mathcal{V}$ that is consistent with the partial order induced by the directed graph G. This order naturally induces an order of the edge set $\mathcal{E}$, where edges $e > e'$ if either $\text{tail}(e) > \text{tail}(e')$ or $\text{tail}(e) = \text{tail}(e')$ and $\text{head}(e) > \text{head}(e')$. WLOG, we assume that $\mathcal{E}_{\mathrm{i}}(j) = \emptyset$ for all source nodes $j \in S$, and $\mathcal{E}_{\mathrm{o}}(\rho) = \emptyset$. We will illustrate in Section III-C how to apply our results on a network with $\mathcal{E}_{\mathrm{i}}(j) \neq \emptyset$ for certain $j \in S$.

The network defined above is used to compute a function, where multiple inputs are generated at the source nodes and the output of the function is demanded by the sink node. The computation units with unbounded computing ability are allocated at all the network nodes. However, the computing capability of the network will be bounded by the network transmission capability. Denote by $\mathcal{B}$ a finite alphabet. We assume that each edge can transmit a symbol in $\mathcal{B}$ reliably for each use.

Denote by $\mathcal{A}$ and $\mathcal{O}$ two finite alphabets. Let $f : \mathcal{A}^s \to \mathcal{O}$ be the *target function*, which is the function to be computed via the network and whose ith input is generated at the ith source node. We may use the network to compute the function multiple times. Suppose that the jth source node consecutively generates k symbols in $\mathcal{A}$ denoted by $x_{1j}, x_{2j}, \ldots, x_{kj}$, and the symbols generated by all the source nodes can be given as a matrix $x = (x_{ij})_{k \times s}$. We denote by x_j the jth column of x, and denote by x^i the ith row of x. In other words, x_j is the vector of the symbols generated at the jth source node, and x^i is the input vector of the ith computation of the function f. Define for $x \in \mathcal{A}^{k \times s}$

$$f^{(k)}(x) = \left(f(x^1), f(x^2), \ldots, f(x^k)\right)^\top.$$

For convenience, we denote by x_J the submatrix of x formed by the columns indexed by $J \subset S$, and denote by x^I the submatrix of x formed by the rows indexed by $I \subset \{1, 2, \ldots, k\}$. We equate $\mathcal{A}^{1 \times s}$ with $\mathcal{A}^s$ in this paper.

For two positive integers n and k, a (n, k) (function-computing) network code over network $\mathcal{N}$ with target function f is defined as follows. Let $x \in \mathcal{A}^{k \times s}$ be the matrix formed by symbols generated at the source nodes. The purpose of the code is to compute $f^{(k)}(x)$ by transmitting at most n symbols in $\mathcal{A}$ on each edge in $\mathcal{E}$. Denote the symbols transmitted on edge e by $g^e(x) \in \mathcal{B}^n$. For a set of edges $E \subset \mathcal{E}$ we define

$$g^E(x) = (g^e(x) | e \in E)$$

where $g^{e_1}(x)$ comes before $g^{e_2}(x)$ whenever $e_1 < e_2$. The (n, k) network code contains the encoding function for each edge e, define:

$$h^e : \begin{cases} \mathcal{A}^k \to \mathcal{B}^n, & \text{if } u \in S; \\ \prod_{e' \in \mathcal{E}_{\mathrm{i}}(\text{tail}(e))} \mathcal{B}^n \to \mathcal{B}^n, & \text{otherwise.} \end{cases}$$

Functions $h^e, e \in \mathcal{E}$ determine the symbols transmitted on the edges. Specifically, if e is an outgoing edge of the ith source node, then

$$g^e(x) = h^e(x_i);$$

if e is an outgoing edge of $u \in \mathcal{V} \setminus (S \cup \{\rho\})$, then

$$g^e(x) = h^e\left(g^{\mathcal{E}_{\mathrm{i}}(u)}(x)\right).$$

The (n, k) network code also contains a decoding function

$$\varphi : \prod_{e' \in \mathcal{E}_{\mathrm{i}}(\rho)} \mathcal{B}^n \to \mathcal{O}^k.$$

Define

$$\psi(x) = \varphi\left(g^{\mathcal{E}_{\mathrm{i}}(\rho)}(x)\right).$$

If the network code *computes* f, i.e., $\psi(x) = f^{(k)}(x)$ for all $x \in \mathcal{A}^{k \times s}$, we then call $\frac{k}{n} \log_{|\mathcal{B}|} |\mathcal{A}|$ an *achievable* computing rate, where we multiply $\frac{k}{n}$ by $\log_{|\mathcal{B}|} |\mathcal{A}|$ in order to normalize the computing rate for target functions with different input alphabets. The *computing capacity* of network $\mathcal{N}$ with respect to a target function f is defined as

$$\mathcal{C}(\mathcal{N}, f) = \sup\left\{\frac{k}{n} \log_{|\mathcal{B}|} |\mathcal{A}| \middle| \frac{k}{n} \log_{|\mathcal{B}|} |\mathcal{A}| \text{ is achievable}\right\}.$$

B. Cut Sets and Special Cases

For two nodes u and v in $\mathcal{V}$, denote the relation $u \to v$ if there exists a directed path from u to v in G. If there is no directed path from u to v, we say u is *separated* from v. Given a set of edges $C \subseteq \mathcal{E}$, I_C is defined to be the set of source nodes which are separated from the sink node ρ if C is deleted from $\mathcal{E}$. Set C is called a *cut set* if $I_C \neq \emptyset$, and the family of all cut sets in network $\mathcal{N}$ is denoted as $\Lambda(\mathcal{N})$. Additionally, we define the set K_C as

$$K_C = \{i \in S | \exists v, t \in \mathcal{V}, i \to v, (v, t) \in C\}.$$

It is reasonable to assume that $u \to \rho$ for all $u \in \mathcal{V}$. Then one can easily see that K_C is the set of source nodes from which there exists a path to the sink node through C. Define $J_C = K_C \backslash I_C$.

The problem also becomes simple when $s = 1$.

Proposition 1. *For a network $\mathcal{N}$ with a single source node and any target function $f : \mathcal{A} \to \mathcal{O}$,*

$$\mathcal{C}(\mathcal{N}, f) = \min_{C \in \Lambda(\mathcal{N})} \frac{|C|}{\log_{|\mathcal{A}|} |f[\mathcal{A}]|},$$

where $f[\mathcal{A}]$ is the image of f on $\mathcal{O}$.

Proof: We first show that the right hand side in the above equation could be achieved. Let

$$M = \min_{C \in \Lambda(\mathcal{N})} |C|.$$

Fix integers k and n such that $k/n \log_{|\mathcal{B}|} |\mathcal{A}| \leq \mathcal{C}(\mathcal{N}, f)$, which implies

$$|f[\mathcal{A}]|^k \leq |\mathcal{B}|^{nM}. \tag{1}$$

By the max-flow min-cut theorem of directed acyclic graphs, there must exist M edge-disjoint paths from the source node to the sink node, and each of them can transmit one symbol in $\mathcal{B}$. We can apply the following network computing code to compute $f^{(k)}$: The source node first computes $f^{(k)}(x)$, and then encodes $f^{(k)}(x)$ into a unique sequence y in $\mathcal{B}^{nM}$ using a one-to-one mapping m, whose existence is guaranteed by (1). The network then forwards y from the source node to the sink node by n uses of the M paths. The sink node decodes by $m^{-1}(y)$.

We then prove the converse. Suppose that we have a (n, k) code with $k/n \log_{|\mathcal{B}|} |\mathcal{A}| > \mathcal{C}(\mathcal{N}, f)$, which implies

$$|f[\mathcal{A}]|^k > |\mathcal{B}|^{nM}. \tag{2}$$

Fix a cut set C with $|C| = M$. It can be shown that

$$\psi(x) = \psi^C(g^C(x)),$$

for certain function ψ^C (see Lemma 5 in Section IV). By (2) and the pigeonhole principle, there must exist $x, x' \in \mathcal{A}^{k \times s}$ such that

$$f^{(k)}(x) \neq f^{(k)}(x'),$$
$$g^C(x) = g^C(x'),$$

where the second equality implies $\psi(x) = \psi(x')$. Thus this network code cannot compute both $f^{(k)}(x)$ and $f^{(k)}(x')$ correctly. The proof is completed. ■

C. Upper Bounds

In this paper, we are interested in the general upper bound on $\mathcal{C}(\mathcal{N}, f)$. The first upper bound is induced by Proposition 1

Proposition 2. *For a network $\mathcal{N}$ with target function f,*

$$\mathcal{C}(\mathcal{N}, f) \leq \min_{C \in \Lambda(\mathcal{N}) : I_C = S} \frac{|C|}{\log_{|\mathcal{A}|} |f[\mathcal{A}^s]|}.$$

Proof: Build a network $\mathcal{N}'$ by joining all the source nodes of $\mathcal{N}$ into a single "super" source node. Since a code for network $\mathcal{N}$ can be naturally converted to a code for network $\mathcal{N}'$ (where the super source node performs the operations of all the source nodes in $\mathcal{N}$), we have

$$\mathcal{C}(\mathcal{N}, f) \leq \mathcal{C}(\mathcal{N}', f).$$

The proof is completed by applying Proposition 1 on $\mathcal{N}'$ and $\Lambda(\mathcal{N}') = \{C \in \Lambda(\mathcal{N}) : I_C = S\}$. ■

The above upper bound only uses the image of function f. We propose an enhanced upper bound by investigating an equivalence relation on the input vectors of f. We will compare this equivalence relation with similar definitions proposed in [2], [3] in the next section.

Definition 1 (Equivalence Class)**.** For any function $f : \mathcal{A}^s \to \mathcal{O}$, any two disjoint index sets $I, J \subseteq S$, and any $a, b \in \mathcal{A}^{|I|}, c \in \mathcal{A}^{|J|}$, we say $a \overset{(c)}{\equiv} b|_{I,J}$ if for every $x, y \in \mathcal{A}^{1 \times s}$, we have $f(x) = f(y)$ whenever $x_I = a$, $y_I = b$, $x_J = y_J = c$ and $x_{S \setminus (I \cup J)} = y_{S \setminus (I \cup J)}$. Two vectors a and b satisfying $a \overset{(c)}{\equiv} b|_{I,J}$ are said to be (I, J, c)*-equivalent. When $J = \emptyset$ in the above definition, we use the convention that c is an empty matrix.*

Note that the equivalence as defined above does not depend on the structure of the network. However, it will soon be clear that with a network, the division of equivalence classes naturally leads to an upper bound of the network function-computing capacity based on cut sets.

For every f, I, J and $c \in \mathcal{A}^{|J|}$, let $W_{I,J,f}^{(c)}$ denote the total number of equivalence classes induced by $\overset{(c)}{\equiv}|_{I,J}$. Given a network $\mathcal{N}$ and a cut set C, let $W_{C,f} = \max_{c \in \mathcal{A}^{|J_C|}} W_{I_C,J_C,f}^{(c)}$. Our main result is stated as following. The proof of the theorem is presented in Section IV).

Theorem 3. *If $\mathcal{N}$ is a network and f is a target function, then*

$$\mathcal{C}(\mathcal{N}, f) \leq \min_{C \in \Lambda(\mathcal{N})} \frac{|C|}{\log_{|\mathcal{A}|} W_{C,f}} := \textit{min-cut}(\mathcal{N}, f).$$

III. DISCUSSION OF UPPER BOUND

In this section, we first give an example to illustrate the upper bound. We compare our result with the existing ones, and proceed by a discussion about the tightness of the bound.

A. An Illustration of the Bound

First we give an example to illustrate our result. Consider the network $\mathcal{N}_1$ in Fig. 1 with the object function $f(x_1, x_2, x_3) = x_1 x_2 + x_3$, where $\mathcal{A} = \mathcal{B} = \mathcal{O} = \{0, 1\}$.

Let us first compare the upper bounds in Theorem 3 and Proposition 2. Let $C_0 = \{e_6, e_7\}$. Here we have

- $|C_0| = 2$, $I_{C_0} = \{3\}$, $J_{C_0} = \{1, 2\}$; and
- For any given inputs of nodes 1 and 2, different inputs from node 3 generate different outputs of f. Therefore $W_{I_{C_0},J_{C_0},f}^{(c)} = 2$ for any $c \in \mathcal{A}^2$ and hence $W_{C_0,f} = 2$.

By Theorem 3, we have

$$\mathcal{C}(\mathcal{N}_1, f) \leq \text{min-cut}(\mathcal{N}_1, f) \leq \frac{|C_0|}{\log_{|\mathcal{A}|} W_{C_0,f}} = 2.$$

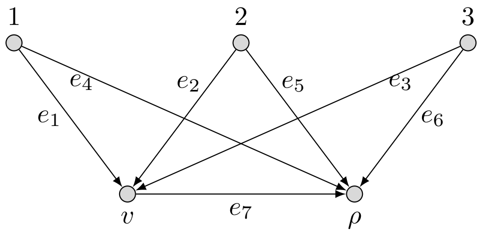

Fig. 1. Network $\mathcal{N}_1$ has three source nodes, $1, 2$ and 3, and one sink node ρ that computes the nonlinear function $f(x_1, x_2, x_3) = x_1x_2 + x_3$, where $\mathcal{A} = \mathcal{B} = \mathcal{O} = \{0, 1\}$.

While Proposition 2 induces that

$$\begin{aligned}\mathcal{C}(\mathcal{N}_1, f) &\le \min_{C\in\Lambda(\mathcal{N}_1):I_C=S} \frac{|C|}{\log_{|\mathcal{A}|}|f[\mathcal{A}^s]|}\\ &= \min_{C\in\Lambda(\mathcal{N}_1):I_C=S} |C|\\ &= 4,\end{aligned}$$

where the first equality follows from $f[\mathcal{A}^s] = \{0, 1\}$, and the second equality follows from

$$\min_{C\in\Lambda(\mathcal{N}_1):I_C=S} |C| = |\{e_4, e_5, e_6, e_7\}| = 4.$$

Therefore, Theorem 3 gives a strictly better upper bound than Proposition 2.

The upper bound in Theorem 3 is actually tight in this case. We claim that there exists a $(1, 2)$ network code that computes f in $\mathcal{N}_1$. Consider an input matrix $x = (x_{ij})_{2\times 3}$. Node i sends x_{1i} to node v and sends x_{2i} to node ρ for $i = 1, 2, 3$ respectively, i.e., for $i = 1, 2, 3$

$$g^{e_i} = x_{1i}, \quad g^{e_{i+3}} = x_{2i}.$$

Node v then computes $f(x^1) = x_{11}x_{12} + x_{13}$ and sends it to node ρ via edge e_7. Node ρ receives $f(x^1)$ from e_7 and computes $f(x^2) = x_{21}x_{22} + x_{23}$ using the symbols received from edges e_4, e_5 and e_6.

B. Comparison with Previous Works

Upper bounds on the computing capacity have been studied in [2], [3] based on a special case of the equivalence class defined in Definition 1. However, we will demonstrate that the bounds therein do not hold for the example we studied in the last subsection.

In Definition 1, when $J = \emptyset$, we will say $a \equiv b|_I$, or a and b are I-equivalent. That is $a \equiv b|_I$ if for every $x, y \in \mathcal{A}^{1\times s}$ with $x_I = a$, $y_I = b$ and $x_{S\setminus I} = y_{S\setminus I}$, we have $f(x) = f(y)$. For target function f and $I \subset S$, denote by $R_{I,f}$ the total number of equivalence classes induced by $\equiv|_I$. For a cut $C \in \Lambda(\mathcal{N})$, let $R_{C,f} = R_{I_C,f}$. For $J \subseteq S$, let $\mathbf{i}_J : J \to \{1, \ldots, |J|\}$ be the one-to-one mapping preserving the order on J, i.e., $\mathbf{i}_J(i) < \mathbf{i}_J(j)$ if and only if $i < j$. Then we have the following lemma:

Lemma 1. *Let I, J be disjoint subsets of S and $J' \subseteq J$. Then for $a, b \in \mathcal{A}^{|I|}$ and $c \in \mathcal{A}^{|J|}$, we have that $a \overset{(c')}{\equiv} b|_{I,J'}$ implies $a \overset{(c)}{\equiv} b|_{I,J}$ where $c' = c_{\mathbf{i}_J(J')}$. In particular, $a \equiv b|_I$ implies $a \overset{(c)}{\equiv} b|_{I,J}$ for all $J \subseteq S \setminus I$ and $c \in \mathcal{A}^{|J|}$.*

Proof: Assume that $a \overset{(c')}{\equiv} b|_{I,J'}$ where $I \subseteq S, J \subseteq S \setminus I, J' \subseteq J$, $a, b \in \mathcal{A}^{|I|}$, $c \in \mathcal{A}^{|J|}$ and $c' = c_{\mathbf{i}_J(J')}$. We want to prove that $a \overset{(c)}{\equiv} b|_{I,J}$.

It suffices to show that $f(x) = f(y)$ for all $x, y \in \mathcal{A}^s$ satisfying $x_I = a, y_I = b, x_J = y_J = c, x_{S\setminus(I\cup J)} = y_{S\setminus(I\cup J)}$. We know that $x_{J'} = (x_J)_{\mathbf{i}_J(J')} = c_{\mathbf{i}_J(J')} = c'$ by definition of the function $\mathbf{i}_J$. Therefore $x_I = a, y_I = b, x_{J'} = y_{J'} = c'$ and $x_{S\setminus(I\cup J')} = y_{S\setminus(I\cup J')}$, which implies $f(x) = f(y)$. The proof is finished. ■

Lemma 2. *Let $I \subset I' \subset S$ and $J = I' \setminus I$. For all $c \in \mathcal{A}^{|J|}$, $a \equiv b|_I$ implies $a' \equiv b'|_{I'}$ where $a'_{\mathbf{i}_{I'}(I)} = a$, $b'_{\mathbf{i}_{I'}(I)} = b$ and $a'_{\mathbf{i}_{I'}(J)} = b'_{\mathbf{i}_{I'}(J)} = c$.*

Proof: We have $a \equiv b|_I$ implies $a \overset{(c)}{\equiv} b|_{I,J}$ by Lemma 1, and $a' \equiv b'|_{I'}$ is equivalent to $a \overset{(c)}{\equiv} b|_{I,J}$ by definition. ■

Lemma 3. *Fix network $\mathcal{N}$ and function f. Then, i) for any $C \in \Lambda(\mathcal{N})$, we have $R_{C,f} \ge W_{C,f}$; ii) for any $C, C' \in \Lambda(\mathcal{N})$ with $C' \subset C$ and $I_{C'} = I_C$, we have $W_{C',f} \ge W_{C,f}$.*

Proof: Let $I = I_C$, $J = J_C$ and $J' = J_{C'}$. Apparently, $J_{C'} \subseteq J_C$. By Lemma 1, for $c \in \mathcal{A}^{|J_C|}$, we have

$$W^{(c')}_{I,J',f} \ge W^{(c)}_{I,J,f},$$

where $c' = c_{\mathbf{i}_J(J')}$. In particular,

$$R_{I,f} \ge W^{(c)}_{I,J,f}.$$

Then $R_{I,f} \ge \max_{c\in\mathcal{A}^{|J|}} W^{(c)}_{I,J,f} = W_{C,f}$.

Fix $c^* \in \mathcal{A}^{|J_C|}$ such that $W^{(c^*)}_{I,J,f} = W_{C,f}$. Then, $W_{C',f} \ge W^{(c'')}_{I,J',f} \ge W^{(c^*)}_{I,J,f} = W_{C,f}$ where $c'' = c^*_{\mathbf{i}_J(J')}$. ■

Define

$$\text{min-cut}_{\text{A}}(\mathcal{N}, f) = \min_{C\in\Lambda(\mathcal{N})} \frac{|C|}{\log_{|\mathcal{A}|} R_{C,f}}.$$

By Lemma 3, we have $\text{min-cut}(\mathcal{N}, f) \ge \text{min-cut}_{\text{A}}(\mathcal{N}, f)$. It is claimed in [2, Theorem II.1] that $\text{min-cut}_{\text{A}}(\mathcal{N}, f)$ is an upper bound on $C(\mathcal{N}, f)$. We find, however, $\text{min-cut}_{\text{A}}(\mathcal{N}, f)$ is not universally an upper bound for the computing capacity. Consider the example in Fig. 1. For cut set $C_1 = \{e_4, e_6, e_7\}$, we have $I_{C_1} = \{1, 3\}$. On the other hand, it can be proved that $R_{C_1,f} = 4$ since i) f is an affine function of x_2 given that x_1 and x_3 are fixed, and ii) it takes 2 bits to represent this affine function over the binary field. Hence

$$\text{min-cut}_{\text{A}}(\mathcal{N}_1, f) \le \frac{|C_1|}{\log_{|\mathcal{A}|} R_{C_1,f}} = \frac{3}{2} < 2 = \mathcal{C}(\mathcal{N}_1, f).$$

For a network $\mathcal{N}$ as defined in Section II-A, we say a subset of nodes $U \subset \mathcal{V}$ is a *cut* if $|U \cap S| > 0$ and $\rho \notin U$. For a cut U, denote by $\mathcal{E}(U)$ the cut set determined by U, i.e.,

$$\mathcal{E}(U) = \{e \in \mathcal{E} : \text{tail}(e) \in V, \text{head}(e) \in \mathcal{V} \setminus U\}.$$

Let

$$\Lambda^*(\mathcal{N}) = \{\mathcal{E}(U) : U \text{ is a cut in } \mathcal{N}\}.$$

Define

$$\text{min-cut}_{\text{K}}(\mathcal{N}, f) = \min_{C \in \Lambda^*(\mathcal{N})} \frac{|C|}{\log_{|\mathcal{A}|} R_{C,f}}.$$

Since $\Lambda^*(\mathcal{N}) \subset \Lambda(\mathcal{N})$, $\text{min-cut}_{\text{K}}(\mathcal{N}, f) \geq \text{min-cut}_{\text{A}}(\mathcal{N}, f)$. It is implied by [3, Lemma 3] that $\text{min-cut}_{\text{K}}(\mathcal{N}, f)$ is an upper bound on $\mathcal{C}(\mathcal{N}, f)$. However, $\text{min-cut}_{\text{K}}(\mathcal{N}, f)$ is also not universally an upper bound for the computing capacity. Consider the example in Fig. 1. For the cut $U_1 = \{1, 3, v\}$, the corresponding cut set $\mathcal{E}(U_1) = C_1 = \{e_4, e_6, e_7\}$. Hence,

$$\text{min-cut}_{\text{K}}(\mathcal{N}_1, f) \leq \frac{|C_1|}{\log_{|\mathcal{A}|} R_{C_1,f}} = \frac{3}{2} < 2 = \mathcal{C}(\mathcal{N}_1, f).$$

Though in general the upper bounds in [2], [3] are not valid, for various special cases discussed in [2], e.g., multi-edge tree networks, these bounds still hold.

C. Tightness

The upper bound in Theorem 3 is tight when the network is a multi-edge tree. We may alternatively prove the same result using [2, Theorem III.3], together with the facts that

$$\text{min-cut}(\mathcal{N}, f) = \min_{C \in \Lambda^*(\mathcal{N})} \frac{|C|}{\log_{|\mathcal{A}|} W_{C,f}}$$

and that for a multi-edge tree network, $W_{C,f} = R_{C,f}$ for $C \in \Lambda^*(\mathcal{N})$.

Theorem 4. *If G is a multi-edge tree, for network $\mathcal{N} = (G, S, \rho)$ and any target function f,*

$$\mathcal{C}(\mathcal{N}, f) = \textit{min-cut}(\mathcal{N}, f).$$

Proof: Fix a pair of (n, k) such that $\frac{k}{n} \log_{|\mathcal{B}|} |\mathcal{A}| \leq \text{min-cut}(\mathcal{N}, f)$. It suffices to show that there exists an (n, k) code computing f on $\mathcal{N}$. We have

$$W^k_{C,f} \leq |\mathcal{B}|^{n|C|} \tag{3}$$

for all $C \in \Lambda(\mathcal{N})$. For node u, define

$$P(u) = \{v \in S | v \to u\},$$
$$\text{prec}(u) = \{v \in \mathcal{V} | (v, u) \in \mathcal{E}\}.$$

For all $u \in G$, $\mathcal{E}_{\text{o}}(u)$ is a cut set, and

$$I_{\mathcal{E}_{\text{o}}(u)} = P(u), \quad J_{\mathcal{E}_{\text{o}}(u)} = \emptyset,$$
$$R^k_{\mathcal{E}_{\text{o}}(u),f} \leq |\mathcal{B}|^{n|\mathcal{E}_{\text{o}}(u)|}. \tag{4}$$

We assign each u a function

$$\gamma_u : \mathcal{A}^{|P(u)|} \to \{1, 2, \ldots, R_{P(u),f}\}$$

such that

$$\gamma_u(x) = \gamma_u(y) \Leftrightarrow x \equiv y|_{P(u)}.$$

For any $s \in \mathcal{A}^{|P(u)|}$, the value $\gamma_u(s)$ determines the equivalence class s lies in. For $x \in \mathcal{A}^{k \times |P(u)|}$, let

$$\gamma^k_u(x) = (\gamma_u(x^1), \gamma_u(x^2), \ldots, \gamma_u(x^k))^\top.$$

Then consider the following (n, k) code, where we claim that $\gamma^k_u(x_{P(u)})$ can be computed by all nodes u, given the initial input $x \in \mathcal{A}^{k \times s}$. This claim is proved inductively with the outline of the code.

Each source node $i \in S$ computes $\gamma^k_i(x_i)$. Note that there are only $R^k_{\{i\},f} \leq |\mathcal{B}|^{n|\mathcal{E}_{\text{o}}(i)|}$ possible outputs of γ^k_i. So we can encode each output into a distinct string of $\mathcal{B}^{n|\mathcal{E}_{\text{o}}(i)|}$, and send the $n|\mathcal{E}_{\text{o}}(i)|$ entries of the string in n uses of the edges in $\mathcal{E}_{\text{o}}(i)$. Thus the claim holds for all source nodes i.

For an intermediate node u, let $m = |\text{prec}(u)|$ and denote $\text{prec}(u) = \{v_1, \ldots, v_m\}$. Assume that the claim holds for all $v \in \text{prec}(u)$. Node u first recovers $\gamma^k_v(x_{P(v)})$ for all $v \in \text{prec}(u)$ using the symbols received from $\mathcal{E}_{\text{i}}(u)$. This is possible by the induction hypothesis and $\mathcal{E}_{\text{o}}(v) \subset \mathcal{E}_{\text{i}}(u)$ for all $v \in \text{prec}(u)$. Then node u fixes $s \in \mathcal{A}^{k \times |P(u)|}$ such that $\gamma^k_{v_j}(s_{\mathbf{i}_{P(u)}(P(v_j))}) = \gamma^k_{v_j}(x_{P(v_j)})$ holds for all $1 \leq j \leq m$, i.e.

$$s^i_{\mathbf{i}_{P(u)}(P(v_j))} \equiv x^i_{P(v_j)}|_{P(v_j)}, \forall 1 \leq j \leq m, 1 \leq i \leq k. \tag{5}$$

Such an s can be found by enumerating the matrices in $\mathcal{A}^{k \times |P(u)|}$.

By (4), we can encode each output of γ^k_u into a distinct string of $\mathcal{B}^{n|\mathcal{E}_{\text{o}}(u)|}$. In this (n, k) code, node u encodes $\gamma^k_u(s)$ and sends the $n|\mathcal{E}_{\text{o}}(u)|$ entries of the string in n uses of the edges in $\mathcal{E}_{\text{o}}(u)$.

We claim that $\gamma^k_u(s) = \gamma^k_u(x_{P(u)})$. As G is a tree, we have

$$P(u) = \cup_{1 \leq j \leq m} P(v_j)$$

where $P(v) \cap P(v') = \emptyset$ if $v \neq v'$. Then define a sequence of strings $\{a_j\}^m_{j=0}$, $a_j \in \mathcal{A}^{1 \times P(u)}$ as follows:

$$a_0 = s^i, a_m = x^i_{P(u)},$$

$$(a_j)_{\mathbf{i}_{P(u)}(P(v_l))} = \begin{cases} s^i_{\mathbf{i}_{P(u)}(P(v_l))}, l > j; \\ x^i_{P(v_l)}, l \leq j. \end{cases}$$

Consider strings a_{j-1} and a_j, for all $1 \leq j \leq m$. In Lemma 2, let

$$\begin{cases} a' = a_{j-1}, \\ b' = a_j, \\ a = (a')_{\mathbf{i}_{P(u)}(P(v_j))} = (s^i)_{\mathbf{i}_{P(u)}(P(v_j))}, \\ b = (b')_{\mathbf{i}_{P(u)}(P(v_j))} = (x^i)_{P(v_j)}, \\ c = (a')_{\mathbf{i}_{P(u)}(P(u) \setminus P(v_j))} = (b')_{\mathbf{i}_{P(u)}(P(u) \setminus P(v_j))}, \end{cases} \tag{6}$$

we have $a' \equiv b'|_{P(u)}$, i.e. $a_{j-1} \equiv a_j|_{P(u)}$ for all $1 \leq j \leq m$. The equivalence then extends to

$$a_0 = s^i \equiv a_m = x^i_{P(u)}|_{P(u)}$$

for all $1 \leq i \leq k$. Therefore $s^i \equiv x^i_{P(u)}|_{P(u)}$ always holds, and we have

$$\gamma^k_u(s) = \gamma^k_u(x_{P(u)}),$$

finishing the induction.

The sink node can compute $\gamma^k_\rho(x)$, identifying $f^{(k)}(x)$. ■

The upper bound in Theorem 3 is not tight for certain cases. Consider the network $\mathcal{N}_2$ in Fig. 2(a) provided in [2]. Note

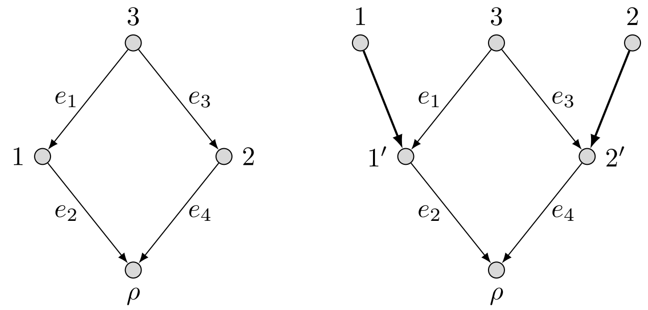

(a) Network $\mathcal{N}_2$ (b) Network $\mathcal{N}_2'$

Fig. 2. Networks $\mathcal{N}_2$ and $\mathcal{N}_2'$ have three binary sources, $\{1, 2, 3\}$ and one sink ρ that computes the arithmetic sum of the source messages, where $\mathcal{A} = \mathcal{B} = \{0, 1\}$. In $\mathcal{N}_2'$, the number of edges from node i to node i' is infinity, $i = 1, 2$.

that in $\mathcal{N}_2$, source nodes 1 and 2 have incoming edges. To match our model described in Section II-A, we can modify $\mathcal{N}_2$ to $\mathcal{N}_2'$ shown in Fig. 2(b), where the number of edges from node i to node i' is infinity, $i = 1, 2$. Every network code in $\mathcal{N}_2'$ naturally induces a network code in $\mathcal{N}_2$ and vise versa. Hence, we have

$$\mathcal{C}(\mathcal{N}_2, f) = \mathcal{C}(\mathcal{N}_2', f).$$

We then evaluate min-cut$(\mathcal{N}_2', f)$. Note that

$$\frac{|C|}{\log_{|\mathcal{A}|} W_{C,f}} < \infty$$

holds only if $|C| < \infty$, and we can thus consider only the finite cut sets. For a finite cut set C, we denote by $C' = C \cap \{e_1, \ldots, e_4\}$. We have $|C'| \leq |C|$ and $J_{C'} \subseteq J_C$, and we claim $I_{C'} = I_C$. Note that $I_{C'} \subseteq I_C$. Suppose that there exists $i \in I_C \setminus I_{C'}$, then there exists a path from i to ρ which is disjoint with C', but shares a subset D of edges with C. Then $D \subset \mathcal{E}_{\mathrm{o}}(i)$ and hence $|D| = 1$. We simply replace the edge in D by an arbitrary edge in $\mathcal{E}_{\mathrm{o}}(i) \setminus C$ and form a new path from i to ρ. This is always possible, since $C \cap \mathcal{E}_{\mathrm{o}}(i)$ is finite while $\mathcal{E}_{\mathrm{o}}(i)$ is not. The newly formed path is disjoint with C, and then we have $i \notin I_C$, a contradiction.

According to Lemma 3, we have $W_{C',f} \geq W_{C,f}$ and hence $\frac{|C'|}{\log_{|\mathcal{A}|} W_{C',f}} \leq \frac{|C|}{\log_{|\mathcal{A}|} W_{C,f}}$. Therefore we can consider only cut sets $C' \subseteq \{e_1, e_2, e_3, e_4\}$. We then have min-cut$(\mathcal{N}_2', f) = 1$, where the minimum is obtained by the cut set $\{e_2, e_4\}$. While for network $\mathcal{N}_2$, it has been proved in [2] that $\mathcal{C}(\mathcal{N}_2, f) = \log_6 4 < 1$. Hence min-cut$(\mathcal{N}_2', f) = 1 > \mathcal{C}(\mathcal{N}_2', f)$.

IV. PROOF OF MAIN THEOREM

To prove Theorem 3, we first give the definition of F-extension and two lemmas.

Definition 2. [F-Extension] Given a network $\mathcal{N}$ and a cut set $C \in \Lambda(\mathcal{N})$, define $D(C) \subseteq \mathcal{E}$ as

$$D(C) = \bigcup_{i \notin I_C} \mathcal{E}_{\mathrm{o}}(i).$$

Then the F-extension of C is defined as

$$F(C) = C \cup D(C).$$

Lemma 4. *For every cut set C, $F(C)$ is a global cut set, i.e.*

$$\forall C \in \Lambda(\mathcal{N}), I_{F(C)} = S.$$

Proof: Clearly, $I_C \subseteq I_{F(C)}$, then it suffices to show that for all $i \notin I_C$, we have $i \in I_{F(C)}$. This is true, since $\mathcal{E}_{\mathrm{o}}(i) \subseteq F(C)$ and $i \in I_{\mathcal{E}_{\mathrm{o}}(i)}$ imply $i \in I_{F(C)}$. ■

Lemma 5. *Consider a (n, k) network code in $\mathcal{N} = (G, S, \rho)$. For any global cut set C, $\psi(x)$ is a function of $g^C(x)$, i.e., $\psi(x) = \psi^C(g^C(x))$ for certain function ψ^C.*

Proof: For a global cut set C of $\mathcal{N}$. Let G_C be the subgraph of G formed by the (largest) connected component of G including ρ after removing C from $\mathcal{E}$. Let S_C be the set of nodes in G_C that do not have incoming edges. Since G_C is also a DAG, S_C is not empty. For each node $u \in S_C$, we have i) u is not a source node in $\mathcal{N}$ since otherwise C would not be a global cut set, and ii) all the incoming edges of u in G are in C since otherwise G_C can be larger. For each node u in G_C but not in S_C, the incoming edges of u are either in G_C or in C, since otherwise the cut set C would not be global. If we can show that for any edge e in G_C, $g^e(x)$ is a function of $g^C(x)$, then $\psi(x) = \varphi(\mathcal{E}_{\mathrm{i}}(\rho))$ is a function of $g^C(x)$.

Suppose that G_C has K nodes. Fix an order on the set of nodes in G_C that is consistent with the partial order induced by G_C, and number these nodes as $u_1 < \ldots < u_K$, where $u_K = \rho$. Denote by $\mathcal{E}_{\mathrm{o}}(u|G_C)$ the set of outgoing edges of u in G_C. We claim that $g^{\mathcal{E}_{\mathrm{o}}(u_i|G_C)}(x)$ is a function of $g^C(x)$ for $i = 1, \ldots, K$, which implies that for any edge e in G_C, $g^e(x)$ is a function of $g^C(x)$. We prove this inductively. First $g^{\mathcal{E}_{\mathrm{o}}(u_1|G_C)}(x)$ is a function of $g^C(x)$ since $u_1 \in S_C$ and hence all the incoming edges of u_1 in G are in C. Assume that the claim holds for the first k nodes in G_C, $k \geq 1$. For u_{k+1}, we have two cases: If $u_{k+1} \in S_C$, the claim holds since all the incoming edges of u_{k+1} in G are in C. If $u_{k+1} \notin S_C$, we know that $\mathcal{E}_{\mathrm{i}}(u_{k+1}) \subset \cup_{i=1}^{k} \mathcal{E}_{\mathrm{o}}(u_i|G_C) \cup C$. By the induction hypothesis, we have that $\mathcal{E}_{\mathrm{o}}(u_{k+1}|G_C)$ is a function of $g^C(x)$. The proof is completed. ■

In the following proof of Theorem 3, it will be handy to extend the equivalence relation for a block of function inputs. For disjoint sets $I, J \in S$ and $c \in \mathcal{A}^{1\times|J|}$ we say $a, b \in \mathcal{A}^{k\times|I|}$ are (I, J, c)-equivalent if for any $x, y \in \mathcal{A}^{k\times s}$ with $x_I = a, y_I = b, x_J = y_J = \left(c^\top, c^\top, \ldots, c^\top\right)^\top$ and $x_{S\setminus I\cup J} = y_{S\setminus I\cup J}$, we have $f^{(k)}(x) = f^{(k)}(y)$. Then for the set $\mathcal{A}^{k\times|I|}$, the number of equivalence classes induced by the equivalence relation is $\left(W_{I,J,f}^{(c)}\right)^k$.

Proof of Theorem 3: Suppose that we have a (n, k) code with

$$\frac{k}{n} \log_{|\mathcal{B}|} |\mathcal{A}| > \text{min-cut}(\mathcal{N}, f). \tag{7}$$

We show that this code cannot compute $f(x)$ correctly for all $x \in \mathcal{A}^{k\times s}$. Denote

$$C^* = \arg\min_{C\in\Lambda(\mathcal{N})} \frac{|C|}{\log_{|\mathcal{A}|} W_{C,f}} \tag{8}$$

and

$$c^* = \arg\max_{c\in\mathcal{A}^{|J_{C^*}|}} W^{(c)}_{I_{C^*},J_{C^*},f}. \tag{9}$$

By (7)-(9), we have

$$\frac{k}{n}\log_{|\mathcal{B}|}|\mathcal{A}| > \frac{|C^*|}{\log_{|\mathcal{A}|} W^{(c^*)}_{I_{C^*},J_{C^*},f}},$$

which leads to

$$|\mathcal{B}|^{|C^*|n} < \left(W^{(c^*)}_{C^*,f}\right)^k. \tag{10}$$

Note that $g^{C^*}(x)$ only depends on $x_{K_{C^*}}$. By (10) and the pigeonhole principle, there exist $a, b \in \mathcal{A}^{k\times|I_{C^*}|}$ such that i) a and b are not (I_{C^*}, J_{C^*}, c^*)-equivalent and ii) $g^{C^*}(x) = g^{C^*}(y)$ for any $x, y \in \mathcal{A}^{k\times s}$ with

$$\begin{cases} x_{I_{C^*}} = a, \quad y_{I_{C^*}} = b, \\ x_{J_{C^*}} = y_{J_{C^*}} = (c^{*\top}, c^{*\top}, \ldots, c^{*\top})^\top, \\ x_{S\setminus K_{C^*}} = y_{S\setminus K_{C^*}}. \end{cases} \tag{11}$$

Fix $x, y \in \mathcal{A}^{k\times s}$ satisfying (11) and $f^{(k)}(x) \neq f^{(k)}(y)$. The existence of such x and y is due to i). Since C^* and $D(C^*)$ are disjoint (see Definition 2) and for any $i \notin I_{C^*}$, $x_i = y_i$, together with ii), we have

$$g^{F(C^*)}(x) = g^{F(C^*)}(y).$$

Thus, applying Lemma 5 we have $\psi(x) = \psi(y)$. Therefore, the code cannot computes both $f^{(k)}(x)$ and $f^{(k)}(y)$ correctly. The proof is completed. ■

REFERENCES

[1] A. Giridhar and P. Kumar, "Computing and communicating functions over sensor networks," *Selected Areas in Communications, IEEE Journal on*, vol. 23, no. 4, pp. 755–764, April 2005.

[2] R. Appuswamy, M. Franceschetti, N. Karamchandani, and K. Zeger, "Network coding for computing: Cut-set bounds," *Information Theory, IEEE Transactions on*, vol. 57, no. 2, pp. 1015–1030, Feb 2011.

[3] H. Kowshik and P. Kumar, "Optimal function computation in directed and undirected graphs," *Information Theory, IEEE Transactions on*, vol. 58, no. 6, pp. 3407–3418, June 2012.

[4] A. Ramamoorthy and M. Langberg, "Communicating the sum of sources over a network," *Selected Areas in Communications, IEEE Journal on*, vol. 31, no. 4, pp. 655–665, April 2013.

[5] R. Ahlswede, N. Cai, S.-Y. R. Li, and R. W. Yeung, "Network information flow," *IEEE Trans. Inform. Theory*, vol. 46, no. 4, pp. 1204–1216, Jul. 2000.

[6] S.-Y. R. Li, R. W. Yeung, and N. Cai, "Linear network coding," *IEEE Trans. Inform. Theory*, vol. 49, no. 2, pp. 371–381, Feb. 2003.

[7] R. Koetter and M. Medard, "An algebraic approach to network coding," *IEEE/ACM Trans. Networking*, vol. 11, no. 5, pp. 782–795, Oct. 2003.

[8] R. Appuswamy and M. Franceschetti, "Computing linear functions by linear coding over networks," *Information Theory, IEEE Transactions on*, vol. 60, no. 1, pp. 422–431, Jan 2014.

[9] B. Rai and B. Dey, "On network coding for sum-networks," *Information Theory, IEEE Transactions on*, vol. 58, no. 1, pp. 50–63, Jan 2012.